

Preface

The past several years have witnessed a great transformative scenario occurring in the computational science by the extremely rapid incursion made by GPU and many-core architecture into the arena of high-performance computing (HPC). At the Supercomputing 2009 meeting in Portland, a group of us (X. Chi, D. A. Yuen, and H. Tufo from University of Colorado) gathered together and formulated plans for a GPU conference to be held in China. First, it was supposed to be held in Shanghai as proposed by Jifeng Yao from Shanghai Supercomputing Center, but because of the EXPO 2010, we had to move it to Harbin, which turns out to be a more delightful locale during the summer. Our timing for this conference was quite propitious, as the platforms armed with GPU and fast networking carried the day and captured three out of the top five slots of the TOP-500 list in Supercomputing 2010 in New Orleans. This auspicious event justified our thinking back in 2009 about the potential importance of GPU computing.

The first International Workshop of GPU Solutions to Multiscale Problems in Science and Engineering (GMP-SMP2010) kicked off July 26, 2010 in Harbin, the capital of Heilongjiang province in Northeast (Dongbei) China. The workshop was organized by the Supercomputing Center of Computer Network Information Center, Chinese Academy of Sciences (CAS) in Beijing, the Graduate University of CAS, and the Chinese Society of Theoretical and Applied Mechanics. Nearly 100 computational experts and scholars from the world's well-known universities and institutes such as the University of Houston, the National Astronomical Observatory of Japan, Hong Kong Baptist University, the University of Chicago, the Tokyo Institute of Technology, Japan, Brown University, the University of Amsterdam, the University of Erlangen-Nuremberg, National Center of Atmospheric Center, University of Minnesota, Macalester College, University of Bonn, the Chinese Academy of Science, Tsinghua University, Peking University, Fudan University, China, University of Science and Technology of China (USTC), attended this international workshop.

The leadoff talk was given by Professor Xuebin Chi, the Director of the Supercomputing Center of CNIC, CAS. He stressed strongly that a revolutionary change is brewing in the field of technologies and applications of high-

performance computing due to the rapid development of GPU and many-core technology. He further expressed the hope that the discussion on GPU applications to multiscale problems in science and engineering be extremely fruitful. Following Professor Chi's speech, Professor David Yuen from University of Minnesota also gave an exciting opening talk, encouraging strongly vigorous participation by the students from China and the international contingent.

During the two-day workshop period, the attendees discussed topics on GPU solutions to multiscale problems in science and engineering. The workshop consisted of three sessions, the keynote lecture session, the invited lecture session, and the student and poster session. The keynote lectures are "*Development and application of a HPC system for multiscale discrete simulation Mole-8.5*" given by Professor Wei Ge from IPE, CAS, "*Acceleration for energy efficient, cost effective HPC*" given by Professor Lennart Johnsson from the University of Houston, USA, "*Practical Random Linear Network Coding on GPUs*" given by Professor Xiaowen Chu from the Hong Kong Baptist University, China and "*GRAPE and GRAPE-DR*" given by Professor Jun Makino from the National Astronomical Observatory of Japan, Japan. In the subsequent talks, issues on seeking GPU solutions to multiscale problems were addressed from different viewpoints, such as focusing on high-performance computing methods and algorithms, efficient software implementation techniques, the construction of scientific computing environment, the mainstream development trends, as well as other GPU-related issues in scientific computing and visualization technology. In the second afternoon students, both graduates and undergraduates from USA and China, gave a bulk of the presentations as well as interesting posters.

In the closing banquet, Professor David Yuen, on behalf of the organizing committee, announced that the conference proceedings will be published soon and the workshop will be held again in 2011, and Lanzhou was mentioned as a potential place. Plans for a book from the workshop were also laid out at the closing banquet on July 28.

As far as we are aware of, the few extant books on GPU and multi-core computing (e.g., Kirk and Hwu 2010; Kurzak et al. 2011) are written by computer scientists. Thus, there is now a dire need for a book with a strong applicational bent in order to encourage more people to join the GPU game.

Our papers are divided into two types: long expository papers with ample illustrations and examples and short contributions with a particular applicational focus in mind. The book is divided into eight sections. In section 1 we begin with this preface. Then we follow with an article by Matt Knepley and David Yuen - which addresses the reasons why scientists and engineers should be considering GPU. This is followed by a chapter by Lang, Wang, and Yuen offering photos of the workshop itself. This book is rather unique at this time has articles, spanning many different areas in science and engineering, which show the health of this burgeoning field. This book is to be contrasted with the recent book edited by Kurzak et al. (2011), which focused mainly on hardware and algorithms and is devoted to a computer science audience rather than people in computational science.

This book covers aspects of hardware and green computing in section 2. We discuss in section 3 software libraries from both China and the USA., namely PARRAY by Chen Yifeng and PETSc by the group at Argonne. In section 4 we shows the industrial applications on GPU. There Wei Ge's group performed an outstanding job in getting a super performance from a GPU-CPU system using thousands of GPU. In section 5, we shows the inroads made by GPU in density-functional theory and electronic structures. Section 6 deals with geophysical and fluid dynamical applications. We see the 145 Tflop performance on a weather code from the Japanese group at Tokyo Institute of Technology. There is also application of GPU on 3D elastic-wave propagation by Taro Okamoto and also by S. Song and his colleagues. In section 7, on algorithms and solvers have a thorough discussion of multigrid solver as applied to industrial problems by H. Koestler from Germany. In the final section on visualization, we present a variety of application on imaging and microtomography using GPU.

We plan to communicate broadly about the potential of GPU and many-core computing to the scientific and engineering community out there and not restricted to computer scientists. We hope this book will give the future perspectives of GPU to scientists and engineers and will stimulate further growth in this field. We thank both the Chinese Academy of Sciences and the OCI program of National Science Foundation for their generous support. We are very grateful for the help provided by Xianyu Lang, Qing Zhao, and Yichen Zhou in preparation of this volume.

Minnesota, USA
Beijing, China

David A. Yuen
Long Wang
Xuebin Chi

References

- Kirk DB, Hwu WW (2010) Programming massively parallel processors: a hands-on-approach. Elsevier, Amsterdam, p 256
- Kurzak J, Bader DA, Dongarra J (eds) (2011) Scientific computing with multicore and accelerators. CRC Press, Boca Raton, p 480

GPU Solutions to Multi-scale Problems in Science and
Engineering

Yuen, D.A.; Wang, L.; Chi, X.; Johnsson, L.; Ge, W.;

Yaolin, S. (Eds.)

2013, XIII, 693 p., Hardcover

ISBN: 978-3-642-16404-0